




Intel Silicon Innovation

Fueling New Solutions for
the Digital Planet



Enabling the Digital Planet

"Moore's Law remains a fundamental enabler of our growth, and it's alive and well at Intel. But the way we and our customers look at Moore's Law has changed. Moore's Law isn't just about more transistors. It's also about how creatively you use those transistors. Intel continues to drive pure performance, generation after generation after generation. We're also using our silicon budget more creatively to enable innovative form factors, integrated wireless and increased ease of use, security and manageability."

Paul S. Otellini
President and Chief Operating Officer
Intel Corporation

As the raw material of the information age, Intel silicon technology has already enabled today's digital planet. Now, as Intel moves deeper into nanotechnology and new capabilities converge onto silicon, the technology's impact promises to be even more exciting. See how Intel is advancing the state-of-the-art in silicon what Intel's breakthroughs mean for the way the world works and plays.

Convergence has happened.

Consumer Electronics. PCs. Telephony. Entertainment.

Increasingly, new products blur the boundaries. PCs make phone calls and show movies with outstanding sound and video quality. Digital phones incorporate video cameras and music players.

Convergence is driving dramatic growth in digital information as entertainment content swells the volumes of digital data. A simple measure: the number of HTML Web pages is doubling every 18 months or so.

The digital universe itself is expanding. Internet use, pegged at 150 million in 1998, is approaching the billion-user mark. With the spread of broadband, these users are joining a "connect anytime, anywhere" society—and marketplace. Innovation is exploding as companies pioneer next-generation applications and services for a converged, connected, digital planet.

Silicon semiconductor technology is the enabler. And Intel is at the forefront.

Silicon is the raw material for the 21st century, just as steel was for the 20th century. And the silicon semiconductor industry—led in part by Intel's technology advances—has delivered a continuing, dramatic spiral of rapid cost reduction and exponential value creation that is unequalled in history.

Because of the cumulative impact of these spiraling increases in capability, silicon—the raw material of the microprocessor—powers the economy and the Internet, runs everything from digital phones and PCs to stock markets and spacecraft, and enables today's information-rich, converged, digital world.

Now, Intel is approaching the billion-transistors-per-processor threshold, giving Intel's product developers a larger canvas on which to paint and make silicon technology the innovation catalyst for a much broader range of industries. If silicon technology's past has been dramatic, its future promises to be even more pervasive and profound. The next several billion users will come largely from Asia, eastern Europe, Latin America, the Middle East and other parts of the world. These users represent the next great growth opportunity for the technology industry, and—more importantly—an opportunity to make profoundly positive changes to the human condition globally. But these changes won't happen without enormous innovation and investment by the technology industry to design next-generation products tailored to the user and environmental needs of the next billion users. Silicon technology advances will pave the way.


Recipe for Success

Accomplish the Near Impossible. Repeat.

In 1965, when the planar integrated circuit (IC) was just four years old and a state-of-the-art IC had maybe 30 transistors, Intel co-founder Gordon E. Moore observed that his engineering team had been able to make transistors smaller and fit roughly twice as many on an IC each year for the past several years. In a now legendary article for Electronics magazine, Moore extrapolated that the doubling could continue for another decade. The doubling cycle was later extended to 24 months, to compensate for expected increases in the complexity of semiconductors, and Moore's prediction has held true ever since—far longer than he ever anticipated.


Fulfilling the predictions of Moore's Law means accomplishing the near impossible—over and over again. Intel's silicon chip designers continue to make transistors smaller and smaller, which means reducing process geometries—scaling (or shrinking) the nominal feature size of the devices populating and powering the silicon chip. Scaling the process geometries makes more space available to bring more transistors, as well as to converge different types of devices and functions, onto the chip.

In the last decade or so, Intel has shrunk its process geometries by an order of magnitude, going from just under 1 micron (approximately 1/100th the width of a human hair) to well into nanotechnology range. In the coming decade, Intel's process geometries will approach the physical limits of atomic structure, bringing new challenges relating to power, heat, and particle behavior. Intel is already manufacturing microprocessors with some features as thin as five atomic layers. Because of shrinking process geometries, today's Intel® Itanium® 2 processor packs 410 million transistors onto a piece of silicon wafer not much larger than your fingernail. Intel has already demonstrated a microprocessor with 1.7 billion transistors.



Moore's Law Means More Performance

Processing power, measured in millions of instructions per second (MIPS), has risen because of increased transistor counts.



Moore's Law Means Decreasing Costs

Packing more transistors into less space has dramatically reduced their cost and the cost of the products they populate.

1. Estimate only

More Transistors, More Power to Innovate

What are the implications of having millions of transistors on a chip? As transistor counts climb, so does the ability to increase device complexity and integrate the convergence of many varieties of capabilities on a chip. This combination of count, complexity, and convergence creates a richer set of resources to make silicon devices even more capable and improve the flexibility and cost-effectiveness with which they can be applied.

Intel is not only putting more transistors on each generation of processors. It is using its rising transistor budgets to provide value-added and in some cases revolutionary new capabilities—from Hyper-Threading Technology and dual-processor cores to integrated communications and security capabilities. As these value-added capabilities become realized in Intel's chips, speed becomes just one way of gauging the increased performance and capacity of each new silicon chip generation.

Here's how it works.

Count. The more Intel's designers can integrate on a chip, the greater the potential for higher performance and new capabilities. Equally important: Smaller transistors reduce the space necessary for a particular logic or memory function, freeing room for new devices.


Complexity. With more transistors available, engineers can use their "transistor budget" more creatively. Within the traditional realms of logic and memory, these new and more complex uses include dual core architectures, which provide greater application flexibility, and the incorporation of air gaps, moving parts, antennae, or other new structures to add functionality through silicon. Gordon Moore has often cited rising complexity as a key factor in realizing Moore's Law.

Convergence. Convergence brings diverse new functions and heterogeneous technologies onto a single chip to maximize integration and functionality. Convergence offers extra opportunity for innovation, raises the value of semiconductor products, and enables traditionally expensive, niche technologies to enter the mainstream by exploiting Intel's silicon technology unmatched price/performance curve. Opportunities for convergence onto silicon chips range from microfluidics and biology to wireless and optical communications.

The Power Challenge

When it comes to power, making transistors smaller and putting more of them into a small space is a good news, bad news story. The good news: Smaller transistors consume less power, and it takes less voltage to drive them. The bad news: Increasing density of ever faster transistors means the overall chip consumes more power and generates more heat. In addition, power leakage becomes more problematic with shrinking feature sizes, wasting a higher portion of total microprocessor power.

Power is a challenge to the entire semiconductor industry and nothing new to Intel. Our leadership in developing solutions to this ongoing challenge often means we are first to explore new territory and encounter new challenges. But, that's where we've always been happiest—and where our technologists do their best work. Intel is exploring, developing and implementing a variety of creative and effective power-saving techniques including new transistor and interconnect structures and materials, innovative approaches to circuit and micro-architecture design, advanced packaging materials, improvements to system components, and software optimization techniques. Intel's 65 nanometer (nm) process generation will combine several advanced power-saving technologies, including lower capacitance interconnects, sleep transistors and second-generation strained silicon—all working together to drive performance higher while reducing power consumption, current leakage and heat.



65 nm Sleep Transistor Reduces Leakage Power

Advancing the State of the Art

Each new silicon process generation that Intel delivers is the result of Intel innovation on a variety of fronts.

- As transistor counts rise, processors consume more power and generate more heat—at an accelerating rate. Finding ways to enhance power efficiency and manage heat are crucial.
- New transistor materials and structures are required to meet new performance and speed objectives.
- Different types of interconnect are necessary to speed signal transmission between devices.
- Lithography—the process of printing the intricate patterns on silicon—must break new barriers as feature sizes become ever smaller.
- Packaging must become more sophisticated to meet ever more stringent thermal management, power delivery, interconnect density and integration requirements.

To maintain the pace of innovation, these goals must be met in a cost-effective manner suitable for high-volume manufacturing.

Breakthroughs in Silicon Na

Focusing on fundamentals of silicon design and manufacture—and combining improvements and innovations in manufacturing process and technology, transistor structure and materials, and silicon packaging—Intel has made dramatic breakthroughs to improve performance, hold the power envelope steady, and extend its silicon technology leadership. Intel researchers are pursuing evolutionary advances to these and similar technologies, as well as investigating the feasibility of revolutionary futuristic approaches.


Manufacturing Process and Technology

Lithography is the process used to print the intricate lines and patterns that define integrated circuits onto silicon wafers. Smaller transistors and finer lines require light with ever shorter wavelengths.

Intel's current 90 nm process technology—the most advanced silicon process in volume production anywhere in the world—prints individual lines smaller than a virus and 1,000 times thinner than a human hair. In a universe where smaller is better, Intel's 90 nm process produces the smallest CMOS transistors in volume production, measuring only 50 nm in gate length. It implements the thinnest gate oxide ever used in production—just 1.2 nm, less than five atomic layers thick. And its seven-layer copper interconnects, one layer thicker than the previous process generation, allow for greater transistor density per cubic centimeter.

Intel's 90 nm process also represents the first high-volume production use of strained silicon, a sophisticated technique Intel developed for speeding up transistors. Intel's strained silicon technique stretches or compresses the grid-like pattern that silicon atoms naturally form, enabling electrons to flow faster with less resistance. While other companies are exploring their own versions of strained silicon, Intel has already developed a breakthrough approach that is cost-efficient and effective for both NMOS and PMOS transistors and avoids increasing defect density—stumbling blocks that had previously made strained silicon impractical. Intel uses its 90 nm process to make over one million microprocessors each week, giving the company and its customers significant advantages in performance, power efficiency and cost.

Intel's 90 nm Transistor



notechnology

65 nm: Further into Nanotechnology

Continuing the march deeper into nanotechnology is Intel's next-generation, 65 nm process. In August, 2004, Intel built fully functional static RAM chips with more than half a billion transistors using this process, and is on track to put the process into production on its state-of-the-art 300 mm wafers in 2005.

Intel's 65 nm process will feature transistors whose gate length is just 35 nm—so tiny that about 100 of these gates could fit inside the diameter of a human red blood cell. Intel's 65 nm process will again set the standard as the industry's smallest and highest-performing CMOS transistors in high-volume production. This process generation will integrate a refinement of Intel's high-performance strained silicon, and will increase density yet again with eight copper interconnect layers. The drive current—a measure of the speed at which the transistors can be turned on and off—is 13 percent faster than the 90 nm generation on NMOS transistors and 16 percent on PMOS transistors. This advance contributes to faster transistors and ultimately higher clock frequencies.

Intel's 65 nm generation will also allow for SRAM cell sizes of $0.57 \mu\text{m}^2$, enabling the integration of larger in-processor caches that enhance performance with minimal increases in power consumption. An innovative transistor design—sleep transistors—will provide the ability to turn off unused blocks of transistors and dramatically reduce active power consumption and current leakage.


EUV Lithography: Tiny Lightwaves

Intel's capability to shrink transistor feature sizes have resulted in dimensions smaller than usable light wavelengths and in doing so have created a looming challenge. Intel has been leading the research effort on Extreme Ultraviolet (EUV) lithography, and expects to use this breakthrough to address it. Utilizing light with a wavelength of 13.4 nm—more than 10 times shorter than normal optical alternatives—EUV lithography uses reflective optics rather than transmissive optics, making it possible to lay down lines far smaller than 50 nm.


Intel has reached several important milestones toward deployment of EUV lithography and expects to use it to build chips on its 32 nm process by the end of the decade. In July 2004, Intel announced it had installed the world's first commercial EUV lithography tool. The micro exposure tool (MET) is not just a stand-alone tool; rather, it's part of Intel's integrated EUV lithography process line—another world first—and is linked with an automated track that includes resist coating and developing operations. Intel plans to develop its EUV masks in house, and has successfully established an EUV mask pilot line.

Advanced Manufacturing Enhances the Impact


To maximize the impact of its process technology leadership, Intel focuses on continuous advances in manufacturing technology. Intel pioneered the industry's move to 300 millimeter (12-inch) wafers on its 130 nm process technology, cutting costs by approximately 30 percent and energy and water requirements by 40 percent per chip. Now, with its 90 nm process in production on these larger wafers, Intel proliferates the benefits of its advanced silicon technology through higher volume, more cost-effective manufacturing.



Intel's Micro Exposure Tool (MET)



$0.57 \mu\text{m}^2$ SRAM Cell
on Intel 65 nm Process




300 mm vs. 200 mm Wafers

50 Quadrillion Nanotransistors...and Rising

Silicon nanotechnology is a key factor in Intel's silicon leadership. Using the definition of nanotechnology as minimum feature sizes of less than 100 nanometers (nm), Intel entered nanotechnology territory in the year 2000 with the launch of the Intel® Pentium® 4 processor, which had transistor gate widths of less than 70 nm. High-volume fabrication of structures smaller than 100 nm now the norm at Intel. Intel produces more than \$20 billion of nanotechnology products every year, making it one of the world's largest nanotech manufacturers.

Intel is pursuing a host of new, enabling nanotechnologies aimed at maintaining its lead in sub 100 nm transistor scaling. One such enabler, atomic layer deposition (ALD), allows for the self-assembly of molecules one mono-atomic layer at a time, based on sophisticated, naturally occurring chemical interactions. Intel is also pursuing research in carbon nanotubes and silicon nanowires, which are built through controlled manipulation of materials at the molecular level and could become the building blocks for future Intel® chip products. Nanowires are thin strands of silicon atoms that can be used to create transistors and/or interconnects in semiconductor circuits. Carbon nanotubes may also serve to channel heat and help solve future power challenges.



Silicon Nanowires

Source: Morales & Lieber, Science 279, 280 (1998)

Transistor Structure and Materials


Intel has made dramatic breakthroughs in developing innovative structures and materials that not only increase performance, but also help reduce power consumption, heat and current leakage.

Innovative Structure: Tri-gate Transistors


The gate is the part of a transistor whose electrical state determines whether the transistor is on or off. Intel has invented a unique, three-dimensional tri-gate transistor that effectively turns a one-lane road into a three-lane highway, tripling the space available for electrical signals to travel but without taking up more area. Intel's design not only delivers higher performance, but because of its unique structure, it provides reduction in current leakage compared to planar transistors. Intel has prototyped the tri-gate transistor in gate lengths of 30 nm and is targeting to use it later this decade.

Materials Breakthroughs: High-k Gate Dielectric + Metal Gate

Intel has scaled the traditional gate insulator (or dielectric) to as thin as 1.2 nm, the equivalent of only five atomic layers. As this thickness decreases, the electric current leakage through the insulator rises, leading to wasted current and needless heat. To obviate the use of ever thinner insulators, Intel has developed a non-silicon high-capacitance (high-k) gate dielectric and combined it with a metal gate material. This combination drastically reduces current leakage and delivers record performance—and every indication is that the two breakthroughs can be integrated into an economical, high-volume manufacturing process. Intel expects to produce them as early as 2007, as part of Intel's 45 nm process generation.



Future High-k Dielectric for
Reduced Gate Leakage



Intel 30 nm Tri-gate Transistor

Nanotechnology-era Packaging

Packaging, too, must meet greater demands in the nanotechnology era. Packaging is what houses the processor die, protects it from contamination, and acts as an interface between the silicon chip and the rest of the computer system. Intel is pioneering packaging designs that allow for higher transistor density and maximize performance through an optimal balance of electrical, thermal and mechanical capabilities.

Thinner, Stackable Multi-chip Packaging

Intel's multi-chip system in package (SIP) technology integrates two or more silicon die wafers which are thinned by grinding the back of the wafer and removing as much as 90 percent of the silicon. This makes it possible to stack multiple chips into a single package, creating an integrated solution with greater functionality and higher performance in the same or less space. Intel has shipped more than 100 million two-die Stacked Chip Scale Packages (SCSPs) with 7 mil chips and a total package height of 1.4 mm. Three-die SCSPs with package heights of 1.2 mm have moved into production and four-die SCSPs have begun to sample. And, continuing to push the envelope, Intel has announced ultra-thin SCSPs that will stack as many as five ultra-thin memory chips on top of one another.

Thermal Innovations


Intel takes a holistic approach to the thermal management problem. Using sophisticated thermal analysis and measurement tools developed at Intel, our packaging engineers partner with our silicon designers to precisely quantify the heat loads of a chip and optimize the packaging to enable easier dissipation of increasing thermal loads. Significant innovations have also been developed in understanding and optimizing the thermal resistance of the paths through which heat is ducted to the environment. Intel has accomplished this through the development of new highly conductive thermal interface materials (TIMs) and through precise process control in packaging to provide repeatable thermal capability. Intel has also partnered with OEMs to develop and demonstrate innovative heat sink solutions for different market segments resulting in increased silicon performance while minimizing the system footprint.

Optical Interconnect Technology

With individual bus speed between a CPU and its supporting chips estimated to increase an order of magnitude within the next 7 to 10 years, today's copper interconnects could become bandwidth-limited. Intel has demonstrated the feasibility of optical chip-to-chip interconnects at data rates over 1 giga-transfers per second, offering the potential for a faster, less expensive alternative to metal-based data buses connecting CPUs to their supporting chips. Intel is currently at the fore-front of development in CMOS-based lasers that are critical to optical chip-to-chip interconnects.



Multiple Die Stacking in Intel Ultra-thin SCSP



Prototype of Intel Optical Interconnect Technology



Intel Drives a New Cycle of Innovation in Convergence

Intel is innovating and investing to bring the value-generation paradigm of Moore's Law to a variety of new areas. Two particularly exciting areas are the convergence of computing with wireless and optical communications capabilities.

Radio Free Intel

Radio Free Intel is Intel's vision of turning today's wireless chipsets into inexpensive silicon-based radios that only take up a little space, but add enormous capability and flexibility. These devices will combine traditional transistor-based logic functionality with new silicon-based wireless elements, to provide a logic chip product with a fully-integrated low-cost silicon CMOS radio. Working toward radio-readiness on every Intel® processor and other Intel® chips, Intel is running test silicon, using standard high volume CMOS processes rather than costly and exotic materials and processes. Intel technologists are also advancing the state-of-the-art on a variety of related fronts:

- Adaptable, cognitive and dynamically reconfigurable radios that automatically select and reconfigure themselves based on a variety of factors such as location, authentication, and application and usage model.
- Microelectromechanical Systems (MEMS) devices that can create microscopic versions of components such as cantilevers, gears and resonators. A MEMS-based radio frequency (RF) antenna will provide critical functionality for integrating radio capability into silicon.
- Seamless networking capability that enables devices to transparently provide "always best" connectivity.
- Smart multiple-antenna systems that help radio transmitters optimize spectrum usage, maximize bandwidth and reduce interference.

Among the potential applications: Wireless sensor networks that can be used for everything from monitoring and adjusting complex oil refinery operations, to assisting the elderly in retaining their independence.

Silicon Photonics

Intel's work in silicon photonics brings optical networking technologies into silicon, extending the cost, size, and performance advantages of Moore's Law to a new and important arena. Intel is driving to make optical communications practical, affordable and ubiquitous not simply for the long-haul markets and wide area networks, but also for smaller networks as well as chip-to-chip and ultimately on-chip interconnects.

In a dramatic advancement reported in the February 22, 2004 issue of the journal *Nature*, Intel researchers announced a breakthrough transistor-like device that can encode data onto a light beam (an optical modulator). Intel's silicon-based optical modulator operates at >1 GHz. That's already more than 50 times faster than the previous research record of about 20 MHz, and even higher rates should be available in the near future by multiplexing the data streams.

Among other potential benefits, Intel's convergence of silicon and optical communications holds the potential for a dramatically faster Internet. As a substitute for copper interconnects, it promises a higher bandwidth yet cost effective way to accelerate bus performance and avoid the potential problems of signal distortion, impedance mismatch, cross-talk, and electromagnetic interference that can arise as copper interconnects reach higher and higher speeds.

Shaping the Digital Planet

“The most exciting thing about new applications going forward will be the surprises I can not predict. The most important things are usually the ones that people within the industry do not see. They tend to develop outside the industry. I do not know. I just wait to be surprised with the next one that comes along.”


Gordon E. Moore

By miniaturizing and connecting electronics, Intel's silicon technology leadership has helped propel the world into the information age. Now, Intel's continued innovation is shaping the digital planet. The break-away power of having hundreds of millions and, soon, billions, of transistors on a chip, along with the new capabilities enabled by added computing elements and converged communications capability, expand the opportunities to bring silicon's ongoing cycle of falling costs, rapid innovation and rising value to new industries.

The potential impact is truly revolutionary. Intel's silicon technology promises to build a bridge between the physical and electronic universes—connecting the world of atoms to the world of bits, photons to electrons, and electrons to radio waves. And, in a world where much of the population has little or no access to digital technology, Intel's continued advances in silicon technology will help bring the next billion users—many from developing nations or those currently experiencing political and economic turmoil—into the digital community. A wealth of new market spaces will emerge as factors of cost, size, and connectivity continue to bring something magical to virtually every human realm and activity. The only limits are human imagination and ingenuity.

Learn more about the technology and research behind Intel's silicon innovation.


Visit www.intel.com/technology/silicon



Intel 90 nm CMOS Transceiver



MEMS-based Intel RF Switches



Intel® Silicon Optical Modulator



For more Information visit www.intel.com/technology

Unless otherwise noted, source for all images is Intel.

Copyright © 2006 Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Leap ahead, the Leap ahead logo, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Printed in the United States. 1205/YVS/HBD/PDF 305868-002US